

# Ingress Point Spreading: A New Primitive for Adaptive Active Network Mapping

Guillermo Baltra, Robert Beverly, Geoffrey G. Xie

Naval Postgraduate School  
{gbaltra, rbeverly, xie}@nps.edu  
March 10, 2014

PAM 2014



# Outline

**1** Background

2 Methodology

3 Results

4 Future Work



# Why knowing the Internet Topology is important:

## Security:

- Better understanding of connectivity richness among ISPs helps to identify **critical infrastructure** and **vulnerabilities**.
- Improved router level maps will enhance Internet **monitoring** and **modeling** capabilities to identify threats and predict cascading impact of various scenarios.

## Networking Research:

- Topology data is essential to create **new protocols**, design **clean-slate architectures**, or examine **Internet evolution** and economics.



# What is the Topology of the Internet?

## Hard to answer:

- Non-stationary and dynamic (in time).
  - Naturally hides information (difficult to observe).
  - Poorly instrumented (not part of original design).
  - Lack of ground truth.
  - Mapping accuracy depends on the number, location, and probing rate of available Vantage Points (VPs).
  - Topological inferences of paths, aliases, and structure can be brittle or lead to false conclusions.
- 
- Recent research, shows that current measurement tools can benefit significantly from an adaptive approach based on probe training and an understanding of network provisioning (Beverly et al, Donnet et al, Spring et al).

# What is the Topology of the Internet?

## Hard to answer:

- Non-stationary and dynamic (in time).
  - Naturally hides information (difficult to observe).
  - Poorly instrumented (not part of original design).
  - Lack of ground truth.
  - Mapping accuracy depends on the number, location, and probing rate of available Vantage Points (VPs).
  - Topological inferences of paths, aliases, and structure can be brittle or lead to false conclusions.
- 
- Recent research, shows that current measurement tools can benefit significantly from an adaptive approach based on probe training and an understanding of network provisioning (Beverly et al, Donnet et al, Spring et al).

# Outline

1 Background

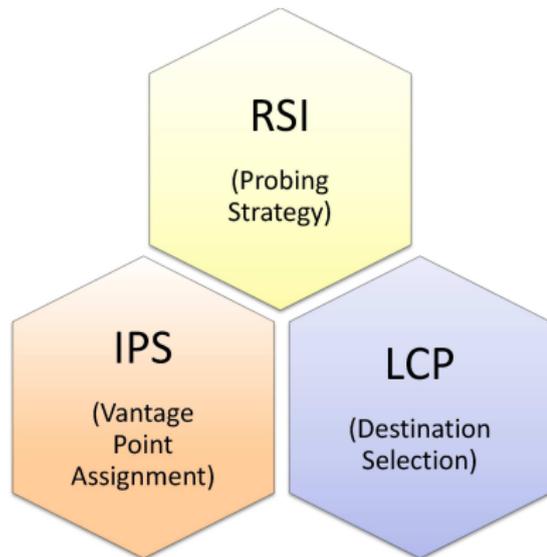
**2 Methodology**

3 Results

4 Future Work



# Probing Strategy



- LCP: Least Common Prefix (Beverly, Berger, Xie [2010])
- RSI: Recursive Subnet Inference
- IPS: Ingress Point Spreading

**Figure:** Three Step Strategy



# Probing Strategy

## Recursive Subnet Inference (RSI)

- Designed to discover the degree of subnetting within networks through an iterative interrogation process.
- Performs a binary search over the target network's address space pruning those branches of the tree that do not reveal new topology information.
- RSI receives as input a network prefix. The address space is divided into 2 halves and probes the center address of each half as defined by the LCP algorithm.
- If a returning probe provides newly discovered interfaces, the procedure is repeated by dividing the corresponding address space into smaller subparts.



# Increasing Probing Efficiency

## Vantage Point Importance

- VPs used in active probing strongly influence the inferred topology (Shavitt, Weinsberg).
- Example 1:
  - CAIDA Ark system, divides the entire routed address space into logical /24 subnetworks.
  - Probes a random address within each /24 using a random VP.
  - Probing every /24 prefix once, constitutes a “cycle.”
  - Assimilates 21 cycles of probing to obtain a high resolution map.



# Increasing Probing Efficiency

## Vantage Point Importance

- For  $N$  cycles and  $M$  VPs, the expected number of unique VPs that explore a given /24 prefix ( $Y$ ) in Ark is given by:

$$E[Y] = M - \frac{(M-1)^N}{M^{N-1}} \quad (1)$$

## Examining one team of CAIDA probing (June, 2013) $M = 18$ VPs:

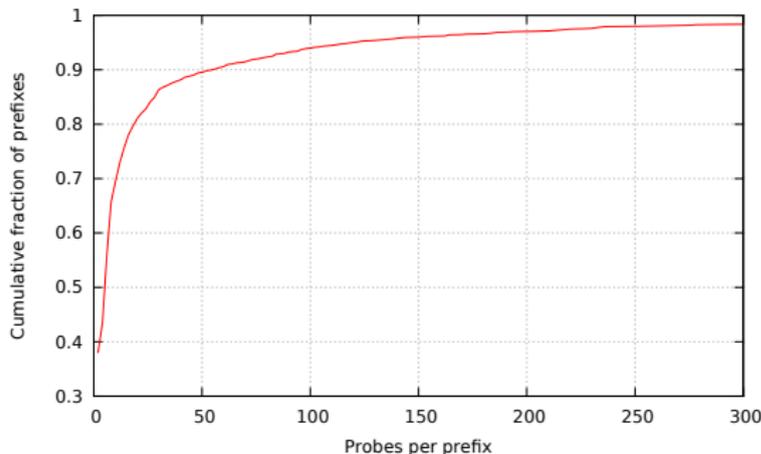
- On average, each /24 in the union of  $N = 21$  cycles is explored by  $E[Y] = 12.6$  VPs.



# Increasing probing efficiency

## Vantage Point Importance

- Example 2: RSI with 60 randomly assigned VPs probing 1500 prefixes selected at random from the global Routeviews BGP tables.

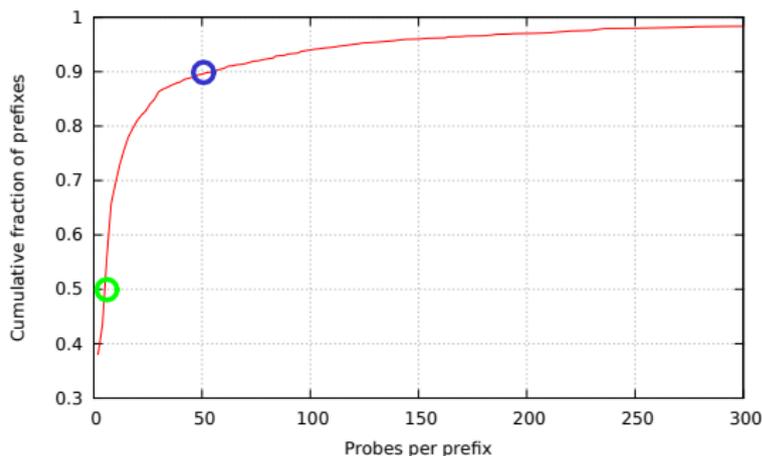


More than half of the prefixes are probed fewer than 10 times, while  $\sim 90\%$  of the prefixes see 50 or fewer probes.

# Increasing probing efficiency

## Vantage Point Importance

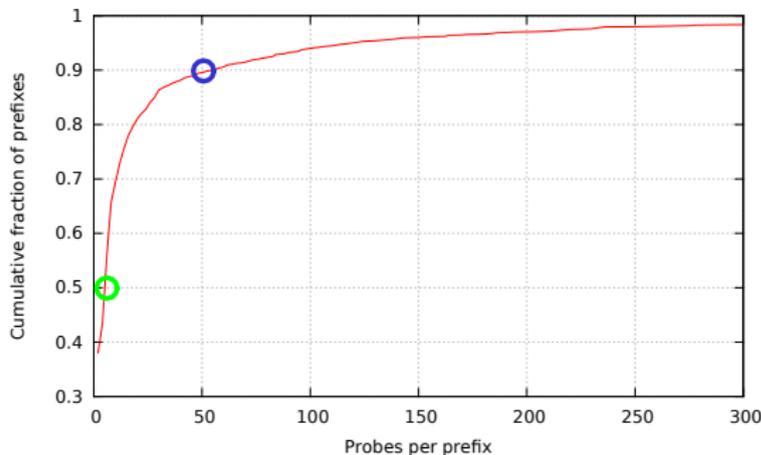
- Example 2: RSI with 60 randomly assigned VPs probing 1500 prefixes selected at random from the global Routeviews BGP tables.



More than half of the prefixes are probed fewer than 10 times, while  $\sim 90\%$  of the prefixes see 50 or fewer probes.

# Increasing probing efficiency

## Vantage Point Importance



More than half of the prefixes are probed fewer than 10 times, while  $\sim 90\%$  of the prefixes see 50 or fewer probes.

- The number of VPs used is frequently less than the total available.
- Even when the number of probes is larger than the number of VPs, using randomly selected VPs is **sub-optimal** (example 1).
- *Therefore, the order in which VPs are employed matters.*

# Increasing probing efficiency

## Ingress Point Spreading (IPS)

- VP selection technique, aimed to discover sources of path diversity into networks.
- Autonomous System (AS) is typically multi-homed and connected with multiple networks.
- IPS infers the number of ingress points for a given network and, then for each new probe, selects the VP with the highest likelihood to traverse a unique ingress point.
- IPS algorithm computes a per-destination network **rank-ordered** list of VPs based on prior rounds of probing.



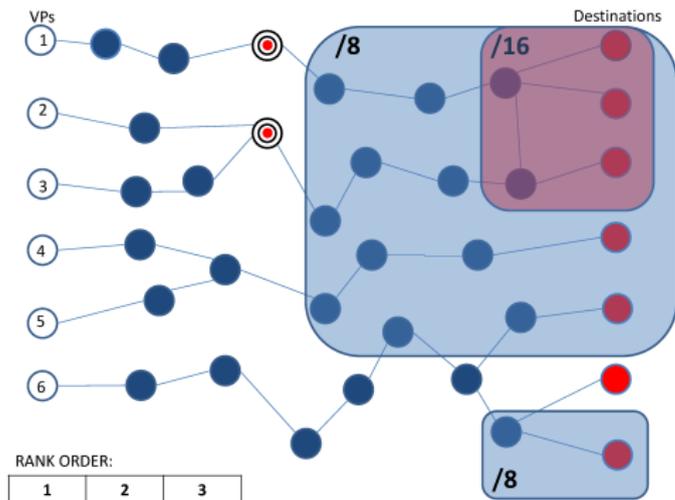
# Ingress Point Spreading

## Notional Prefix

- An expansion to a larger prefix aggregate containing the target prefix.
- By expanding the size of the notional prefix, all VPs can be rank-ordered in order to ensure path diversity.
- *Notional prefix ingress* is the first router interface hop that leads to a next hop whose IP is within the notional prefix.
- Note: *Notional prefix* does not imply relationship to real-world BGP route aggregation.



# Ingress Point Spreading

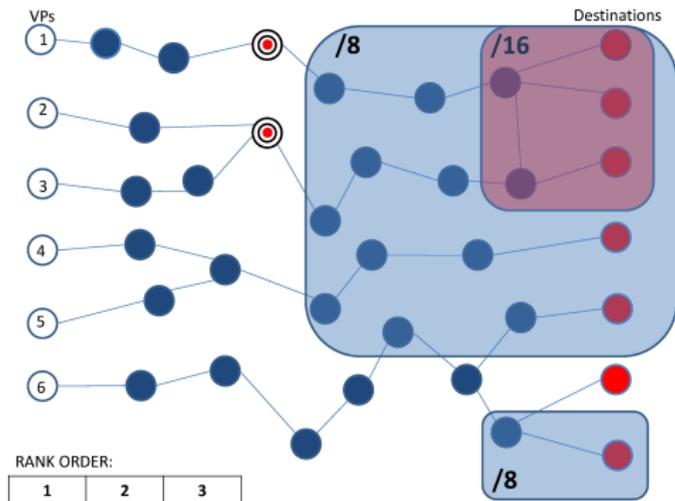


e.g.

- 205.155.0.0/16 is the target prefix (red box).
- /8 is a notional prefix (blue box).
- 6 VPs used.
- Blue circles are hops.
- Red circles are destinations.
- Bullseyes are notional ingress routers.



# Ingress Point Spreading

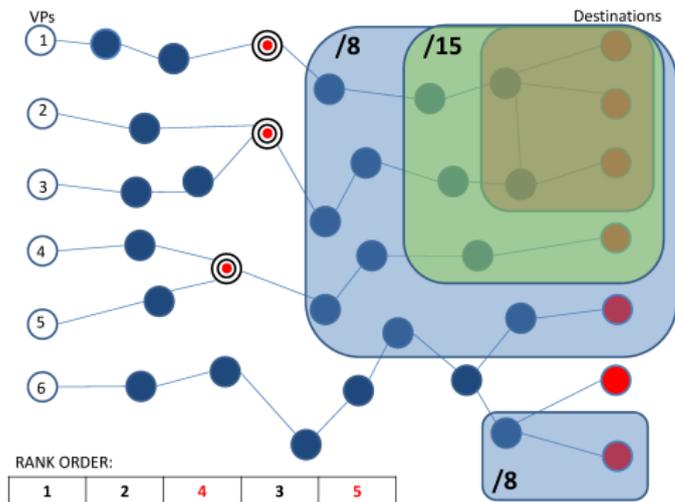


e.g.

- VPs 1 and 2 are selected as the first two VPs in the rank order list, (different ingresses into notional /8 prefix).
- Since VPs 2 and 3 share the same ingress router, the latter is included at the end of the list.
- However, we wish to obtain a total order over all of the VPs.



# Ingress Point Spreading

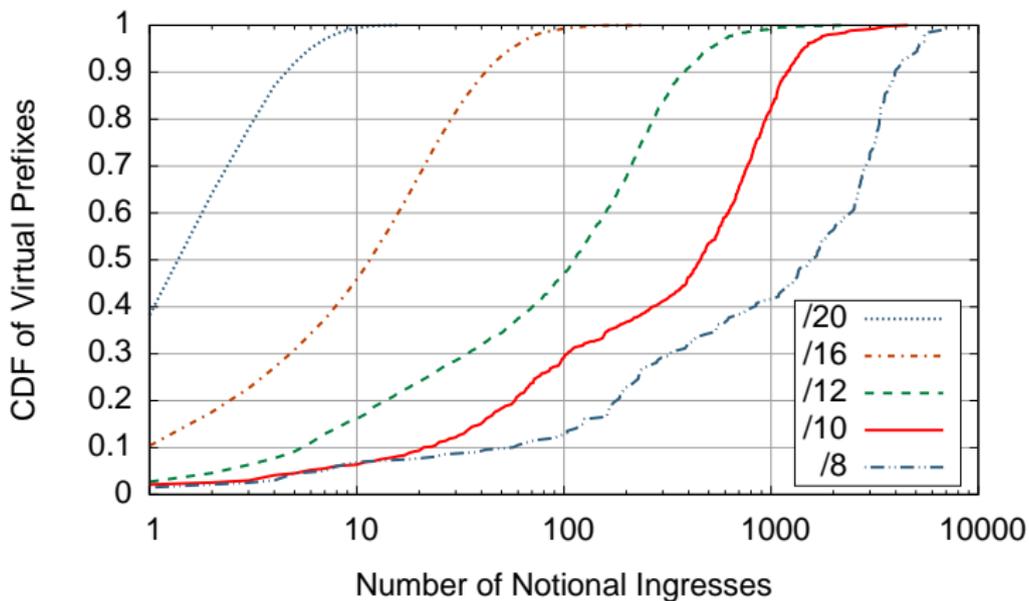


e.g.

- Ingress search space expansion to include  $205.154.0.0/15$  (green box).
- VP 4 becomes the third in the rank-order and VP 5 is included at the end of the list.
- Expansion continues until all VPs are ordered.
- i.e.  $205.152.0.0/14$ ,  $205.152.0.0/13$ , ...,  $205.0.0.0/8$ .

# Notional Prefix

**Figure:** Distribution of Ingresses into Prefixes of Different Logical Size



Data from CAIDA's Ark, June 2-4, 2013.



# Outline

1 Background

2 Methodology

**3 Results**

4 Future Work



# Strategy Evaluation

## IPS compared to popular mapping system, such as Ark:

- Direct comparison with published Ark data is not possible as IPS does not use “teams” of VPs.
- Emulate Ark’s methodology using the same number of VPs for both strategies.
- Pre-probing process: provide IPS with one day’s worth of CAIDA’s topology data (Aug 28, 2013), which demonstrates that IPS is not limited to our own pre-probed data.
- Using IPS and Ark’s strategy,  $\sim 49k$  randomly selected prefixes were probed from 59 globally distributed VPs.



# Strategy Evaluation

Metric	Ark	IPS (Aug. 2013 trained)	IPS (Dec. 2013 trained)
Prefixes Probed	48,905	48,905	48,905
Vertices	464,544	521,513	520,903
Edges	906,680	1,024,295	1,034,101
Probes	4,041,289	2,056,562	2,052,842
Vertices (inside dest)	121,137	135,209	134,575
Vertices (intersection w/ ark)		309,997	309,971
Ingresses	31,138	38,532	39,020
Time	26h 55m	13h 38m	14h 47m

## IPS is significantly more efficient:

- Using ~ 50% the number of probes.
- Taking approximately half the time.
- IPS discovers 211,516 vertices not in Ark.
- Ark discovers 154,547 vertices that IPS does not.

# Strategy Evaluation

Metric	Ark	IPS (Aug. 2013 trained)	IPS (Dec. 2013 trained)
Prefixes Probed	48,905	48,905	48,905
Vertices	464,544	521,513	520,903
Edges	906,680	1,024,295	1,034,101
Probes	4,041,289	2,056,562	2,052,842
Vertices (inside dest)	121,137	135,209	134,575
Vertices (intersection w/ ark)		309,997	309,971
Ingresses	31,138	38,532	39,020
Time	26h 55m	13h 38m	14h 47m

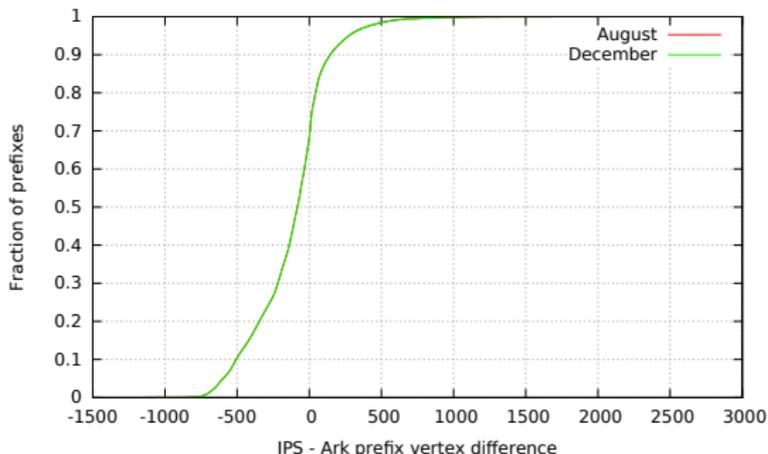
## In terms of performance of IPS against Ark:

- Top 3 prefixes are national ISP networks with hundreds of peering links.
- Bottom 3 prefixes belong to enterprise networks that have small number of peering links.



# Vertex Difference

CDF of per-prefix coverage difference: *IPS* – *Ark*

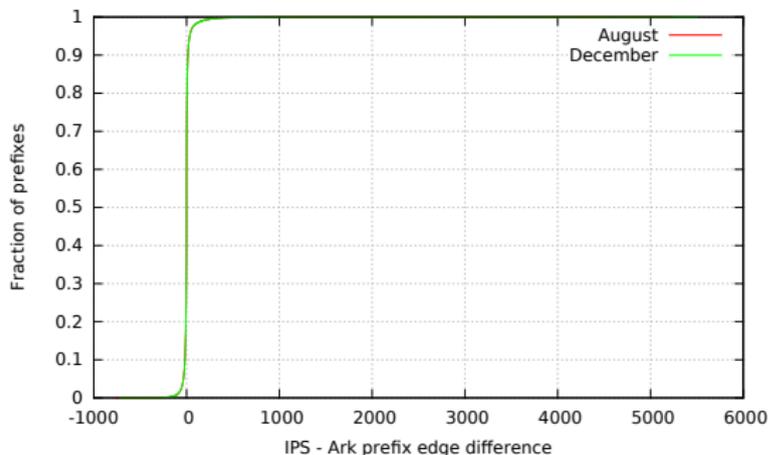


- IPS performs worse than Ark for  $\sim 66\%$  of the prefixes.
- IPS is significantly superior to Ark for a small number of prefixes, thereby contributing to the overall superior topological coverage.



# Edge Difference

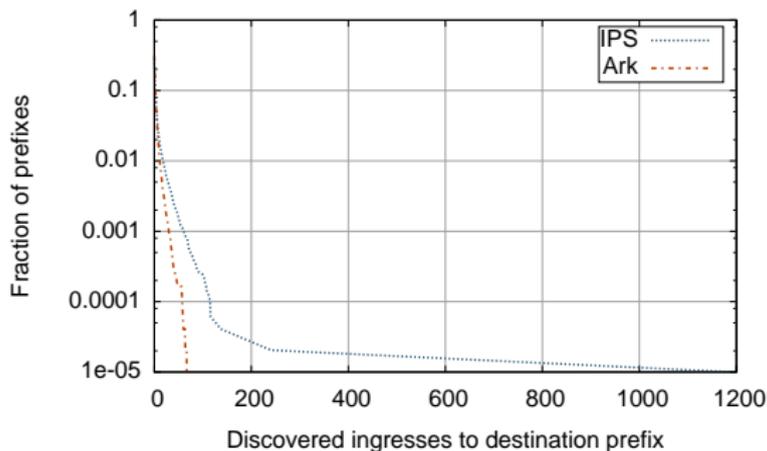
CDF of per-prefix coverage difference: *IPS* – *Ark*



- Similar performance for  $\sim 80\%$  of the prefixes.
- The long tail in the distribution shows that IPS discovers in a small number of prefixes, significantly more topological information.

The fact that IPS performs better on some prefixes while Ark does better on others explains why a high number of interfaces and edges are uniquely discovered by each method.

# Ingress Discovery



- Among destinations where probing within the target network is feasible, IPS finds significantly more ingresses than Ark.
- Neither Ark nor IPS discovers any ingresses for  $\sim 70\%$  of the prefixes (ICMP blocking and other forms of packet filtering).

# Outline

1 Background

2 Methodology

3 Results

**4 Future Work**



# Future Work

**While we have demonstrated promising results by utilizing ingresses to our advantage, significant future work remains:**

- Scale probing by one more order of magnitude to encompass all advertised prefixes on the Internet, and run continually.
- Practical experience has shown that VPs are unreliable, yet IPS cannot simply use the next VP in the ordered list when the preferred VP is down, as the complete ordering is perturbed.
- Some prefixes with significant topology have gone undiscovered by RSI due to the particular deterministic selection of destinations causing early termination.



# Questions

Thanks! Questions?

